arXiv:1605.08386v1 [math.CO] 26 May 2016

HEAT-BATH RANDOM WALKS WITH MARKOV BASES

CAPRICE STANLEY AND TOBIAS WINDISCH

Abstract. Graphs on lattice points are studied whose edges come from a finite set of allowed moves of arbitrary length. We show that the diameter of these graphs on fibers of a fixed integer matrix can be bounded from above by a constant. We then study the mixing behaviour of heat-bath random walks on these graphs. We also state explicit conditions on the set of moves so that the heat-bath random walk, a generalization of the Glauber dynamics, is an expander in fixed dimension.

Contents

1. Introduction 1
2. Graphs and statistics 3
3. Bounds on the diameter 4
4. Heat-bath random walks 8
5. Augmenting Markov bases 14
References 19

1. Introduction

A *fiber graph* is a graph on the finitely many lattice points $\mathcal{F} \subset \mathbb{Z}^d$ of a polytope where two lattice points are connected by an edge if their difference lies in a finite set of allowed moves $\mathcal{M} \subset \mathbb{Z}^d$. The implicit structure of these graphs makes them a useful tool to explore the set of lattice points randomly: At the current lattice point $u \in \mathcal{F}$, an element $m \in \pm\mathcal{M}$ is sampled and the random walk moves along m if $u + m \in \mathcal{F}$ and stays at u otherwise. The corresponding Markov chain is irreducible if the underlying fiber graph is connected and the set $\mathcal{M}$ is called a *Markov basis* for $\mathcal{F}$ in this case. This paper investigates the *heat-bath* version of this random walk: At the current lattice point $u \in \mathcal{F}$, we sample $m \in \mathcal{M}$ and move to a random element in the integer ray $(u + \mathbb{Z} \cdot m) \cap \mathcal{F}$. The authors of [6] discovered that this random walk can be seen as a discrete version of the *hit-and-run* algorithm [15, 26, 16] that has been used frequently to sample from all the points of a polytope – not only from its lattice points. The popularity of the continuous version of the hit-and-run algorithm has not spread to its discrete analog, and not much is known about its mixing behaviour. One reason is that it is already challenging to guarantee that all points in the underlying set $\mathcal{F}$ can be

Date: May 27, 2016.
2010 *Mathematics Subject Classification.* Primary: 05C81, Secondary: 37A25, 11P21.
Key words and phrases. Heat-bath random walks, sampling, lattice points, Markov bases.

reached by a random walk that uses moves from $\mathcal{M}$, whereas for the continuous version, a random sampling from the unit sphere suffices. However, in many situations where a Markov basis is known, the heat-bath random walk is evidently fast. For instance, it was shown in [3] that the heat-bath random walk on contingency tables mixes rapidly when the number of columns is fixed. To work around the connectedness issue, a *discrete hit-and-run* algorithm was introduced in [1] for arbitrary finite sets $\mathcal{F} \subset \mathbb{Z}^d$. At each step in this random walk, a subordinate and unrestricted random walk starts at the current lattice point $u \in \mathcal{F}$ and uses the unit vectors to collect a set of proposals $S \subset \mathbb{Z}^d$. The random walk then moves from u to a random point in $S \cap \mathcal{F}$.

Random walks of the heat-bath type, such as the one presented above, have been studied recently in [8] in a more general context. In this paper, we explore the mixing behaviour of heat-bath random walks on the lattice points of polytopes with Markov bases. Throughout, we assume that a Markov basis has been found already and refer to the relevant literature for their computation [24, 25, 11, 17, 10, 21]. We call the underlying graph of the heat-bath random walk a *compressed fiber graph* (Definition 2.5) and determine in Section 3 bounds on its graph-diameter. We prove that for any $A \in \mathbb{Z}^{m\times d}$ with $\ker_{\mathbb{Z}}(A)\cap\mathbb{N}^d = \{0\}$, the diameter of compressed fiber graphs on $\{u \in \mathbb{N}^d : Au = b\}$ that use a fixed Markov bases $\mathcal{M} \subset \ker_{\mathbb{Z}}(A)$ is bounded from above by a constant as b varies (Theorem 3.15). In contrast, we show that the diameter of conventional fiber graphs grow linearly under a dilation of the underlying polytope (Remark 3.9). This gives rise to slow mixing results for conventional fiber walks as observed in [27]. In Section 4, we study in more detail the combinatorial and analytical structure of the transition matrices of heat-bath random walks on lattice points and prove upper and lower bounds on their second largest eigenvalues. We also discuss how the distribution on the moves $\mathcal{M}$ affects the speed of convergence (Example 4.21). Theorem 5.8 establishes with the *canonical path approach* from [23] an upper bound on the second largest eigenvalue when the Markov basis is *augmenting* (Definition 5.1) and the stationary distribution is uniform. From that, we conclude fast mixing results for random walks on lattice points in fixed dimension.

Acknowledgements. CS was partially supported by the US National Science Foundation (DMS 0954865). TW gratefully acknowledges the support received from the German National Academic Foundation.

Conventions and Notation. The natural numbers are $\mathbb{N} := \{0, 1, 2, \ldots\}$ and for any $N \in \mathbb{N}$, $\mathbb{N}_{>N} := \{n \in \mathbb{N} : n > N\}$ and $\mathbb{N}_{\geq N} := \{N\} \cup \mathbb{N}_{>N}$. For $n \in \mathbb{N}_{>0}$, let $[n] := \{1, \ldots, n\}$. Let $\mathcal{M} \subset \mathbb{Q}^d$ be a finite set, then $\mathbb{Z}\cdot\mathcal{M} := \{\lambda m : m \in \mathcal{M}, \lambda \in \mathbb{Z}\}$ and $\mathbb{N}\mathcal{M}$ is the affine semigroup in $\mathbb{Z}^d$ generated by $\mathcal{M}$. For an integer matrix $A \in \mathbb{Z}^{m\times d}$ with columns $a_1, \ldots, a_d \in \mathbb{Z}^m$, we write $\mathbb{N}A := \mathbb{N}\{a_1, \ldots, a_d\}$. A graph is always undirected and can have multiple loops. The distance of two nodes u, v which are contained in the same connected component of a graph G, i.e. the number of edges in a shortest path between u and v in G, is denoted by $\mathrm{dist}_G(u, v)$. We set $\mathrm{dist}_G(u, v) := \infty$ if u and v are disconnected. A mass function on a finite set Ω is a map $f : \Omega \to [0, 1]$ such that $\sum_{\omega\in\Omega} f(\omega) = 1$. A mass function f on Ω is *positive* if $f(\omega) > 0$ for all $\omega \in \Omega$. A set $\mathcal{F} \subset \mathbb{Z}^d$ is *normal* if it there exists a polytope $\mathcal{P} \subset \mathbb{Q}^d$ such that $\mathcal{P} \cap \mathbb{Z}^d = \mathcal{F}$.

2. Graphs and statistics

We first introduce the statistical framework in which this paper lives and recall important aspects of the interplay between graphs and statistics. A *random walk* on a graph $G = (V, E)$ is a map $\mathcal{H} : V \times V \to [0, 1]$ such that for all $v \in V$, $\sum_{u \in V} \mathcal{H}(v, u) = 1$ and such that $\mathcal{H}(v, u) = 0$ if $\{v, u\} \notin E$. When there is no ambiguity, we represent a random walk as an $|V| \times |V|$-matrix, for example when it is clear how the elements of V are ordered. Fix a random walk $\mathcal{H}$ on G. Then $\mathcal{H}$ is *irreducible* if for all $v, u \in V$ there exists $t \in \mathbb{N}$ such that $\mathcal{H}^t(v, u) > 0$. The random walk $\mathcal{H}$ is *reversible* if there exists a mass function $\mu : V \to [0, 1]$ such that $\mu(u) \cdot \mathcal{H}(u, v) = \mu(v) \cdot \mathcal{H}(v, u)$ for all $u, v \in V$ and *symmetric* if $\mathcal{H}$ is a symmetric map. A mass function $\pi : V \to [0, 1]$ is a *stationary distribution* of $\mathcal{H}$ if $\pi \circ \mathcal{H} = \pi$. For symmetric random walks, the uniform distribution on V is always a stationary distribution. If $|V| = n$, then we denote the eigenvalues of $\mathcal{H}$ by $1 = \lambda_1(\mathcal{H}) \geq \lambda_2(\mathcal{H}) \geq \cdots \geq \lambda_n(\mathcal{H}) \geq -1$ and we write $\lambda(\mathcal{H}) := \max\{\lambda_2(\mathcal{H}), -\lambda_n(\mathcal{H})\}$ for the *second largest eigenvalue modulus* of $\mathcal{H}$. Any irreducible random walk has a unique stationary distribution [14, Corollary 1.17] and $\lambda(\mathcal{H}) \in [0, 1]$ measures the convergence rate: the smaller $\lambda(\mathcal{H})$, the faster the convergence.

The aim of this paper is to study random walks on lattice points that use a set of moves. Typically, this is achieved by constructing a graph on the set of lattice points as follows (compare to [7, Section 1.3] and [24, Chapter 5]).

Definition 2.1. Let $\mathcal{F} \subset \mathbb{Z}^d$ be a finite set and $\mathcal{M} \subset \mathbb{Z}^d$. The graph $\mathcal{F}(\mathcal{M})$ is the graph on $\mathcal{F}$ where two nodes $u, v \in \mathcal{F}$ are adjacent if $u - v \in \mathcal{M}$ or $v - u \in \mathcal{M}$.

A normal set $\mathcal{F} \subset \mathbb{Z}^d$ is finite and satisfies $\mathcal{F} = \text{conv}_{\mathbb{Q}}(\mathcal{F}) \cap \mathbb{Z}^d$. A canonical class of normal sets that arise in many applications, is given by the fibers of an integer matrix:

Definition 2.2. Let $A \in \mathbb{Z}^{m \times d}$ and $b \in \mathbb{N}A$. The set $\mathcal{F}_{A,b} := \{u \in \mathbb{N}^d : Au = b\}$ is the *b-fiber* of A. The collection of all fibers of A is $\mathcal{P}_A := \{\mathcal{F}_{A,b} : b \in \mathbb{N}A\}$. For $\mathcal{M} \subset \ker_{\mathbb{Z}}(A)$, the graph $\mathcal{F}_{A,b}(\mathcal{M})$ is a *fiber graph*.

Let $\mathcal{F}, \mathcal{M} \subset \mathbb{Z}^d$ be finite. If the membership in $\mathcal{F}$ can be verified efficiently – for instance when $\mathcal{F}$ is given implicitly by linear equations and inequalities – then it is possible to explore $\mathcal{F}$ randomly using $\mathcal{M}$ as follows: At a given node $v \in \mathcal{F}$, a uniform element $m \in \mathcal{M}$ is selected. If $v + m \in \mathcal{M}$, then the random walk moves along m to $v + m$ and if $v + m \notin \mathcal{M}$, the we stay at v. Formally, we obtain the following random walk.

Definition 2.3. Let $\mathcal{F} \subset \mathbb{Z}^d$ and $\mathcal{M} \subset \mathbb{Z}^d$ be two finite sets. The *simple walk* is the random walk on $\mathcal{F}(\mathcal{M})$ where the probability to traverse between to adjacent nodes u and v is $|\pm\mathcal{M}|^{-1}$ and the probability to stay at a node u is $|\{m \in \pm\mathcal{M} : u + m \notin \mathcal{F}\}| \cdot |\pm\mathcal{M}|^{-1}$.

The simple walk is symmetric and hence the uniform distribution is a stationary distribution (see also [27, Section 2]). To ensure convergence, the random walk has to be irreducible, that is, the underlying graph has to be connected. The following definition is a slight adaption of the generalized Markov basis as defined in [21, Definition 1].

Definition 2.4. Let $\mathcal{P}$ be a collection of finite subsets of $\mathbb{Z}^d$. A finite set $\mathcal{M} \subset \mathbb{Z}^d$ is a *Markov basis* of $\mathcal{P}$, if for all $\mathcal{F} \in \mathcal{P}$, $\mathcal{F}(\mathcal{M})$ is a connected graph.

We refer to [6, Theorem 3.1] for a proof that for collections $\mathcal{P}_A$, a finite Markov basis always exists and can be computed with tools from commutative algebra (see also [11] for more on the computation of Markov bases). We now introduce a construction of graphs on lattice points that also give rise to implementable random walks, but whose edges have far more reach.

Definition 2.5. Let $\mathcal{F} \subset \mathbb{Z}^d$ and $\mathcal{M} \subset \mathbb{Z}^d$ be finite sets. The *compression* of the graph $\mathcal{F}(\mathcal{M})$ is the graph $\mathcal{F}^c(\mathcal{M}) := \mathcal{F}(\mathbb{Z} \cdot \mathcal{M})$.

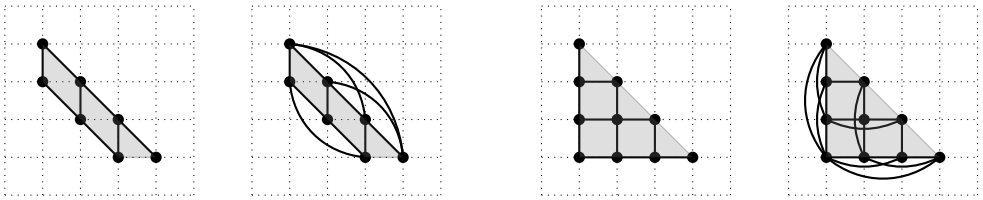

FIGURE 1. Compressing graphs.

Compressing a graph $\mathcal{F}(\mathcal{M})$ preserves its connectedness: $\mathcal{F}(\mathcal{M})$ is connected if and only if $\mathcal{F}^c(\mathcal{M})$ is connected.

3. Bounds on the diameter

In general knowledge of the diameter of the graph underlying a Markov chain can provide information about the mixing time. For random walks on fiber graphs, the chains which we consider, the underlying graph coincides with the fiber graph. In this section, we determine lower and upper bounds on the diameter of fiber graphs and their compressed counterparts. For a finite set $\mathcal{M} \subset \mathbb{Z}^d$ and any norm $\|\cdot\|$ on $\mathbb{R}^d$, let $\|\mathcal{M}\| := \max_{m \in \mathcal{M}} \|m\|$.

Lemma 3.1. *Let $\mathcal{F} \subset \mathbb{Z}^d$ and $\mathcal{M} \subset \mathbb{Z}^d$ be finite sets, then*

$$\operatorname{diam}(\mathcal{F}(\mathcal{M})) \geq \frac{1}{\|\mathcal{M}\|} \cdot \max\{\|u - v\| : u, v \in \mathcal{F}\}.$$

Proof. If $\mathcal{F}(\mathcal{M})$ is not connected, then the statement holds trivially, so assume that $\mathcal{M}$ is a Markov basis for $\mathcal{F}$. Let $u', v' \in \mathcal{F}$ such that $\|u' - v'\| = \max\{\|u - v\| : u, v \in \mathcal{F}\}$ and let $m_1, \ldots, m_r \in \mathcal{M}$ so that $u' = v' + \sum_{i=1}^{r} m_i$ is a path of minimal length, then $\|u' - v'\| \leq r \cdot \|\mathcal{M}\|$ and the claim follows from $\operatorname{diam}(\mathcal{F}(\mathcal{M})) \geq \operatorname{dist}_{\mathcal{F}(\mathcal{M})}(u', v') = r$. □

Remark 3.2. Let $\mathcal{F} \subset \mathbb{Z}^d$ be a normal set. For all $l \in \{-1, 0, 1\}^d$ and $u, v \in \mathcal{F}$ we have $(u - v)^T l \leq \|u - v\|_1$ and thus $\operatorname{width}_l(\mathcal{F}) := \max\{(u - v)^T l : u, v \in \mathcal{F}\} \leq \max\{\|u - v\|_1 : u, v \in \mathcal{F}\}$. Suppose that $u', v' \in \mathcal{F}$ are such that $\|u' - v'\|_1 = \max\{\|u - v\|_1 : u, v \in \mathcal{F}\}$ and let $l'_i := \operatorname{sign}(u'_i - v'_i)$ for $i \in [d]$, then

$$\|u' - v'\|_1 = (u' - v')^T \cdot l' \leq \operatorname{width}_{l'}(\mathcal{F}) \leq \max\{\|u - v\|_1 : u, v \in \mathcal{F}\} = \|u' - v'\|_1.$$

The *lattice width* of $\mathcal{F}$ is $\operatorname{width}(\mathcal{F}) := \min_{l \in \mathbb{Z}^d} \operatorname{width}_l(\mathcal{F})$ and thus Lemma 3.1 gives

$$\|\mathcal{M}\|_1 \cdot \operatorname{diam}(\mathcal{F}(\mathcal{M})) \geq \operatorname{width}(\mathcal{F}).$$

Definition 3.3. Let $\mathcal{P}$ be a collection of finite subsets of $\mathbb{Z}^d$. A finite set $\mathcal{M} \subset \mathbb{Z}^d$ is *norm-like* for $\mathcal{P}$ if there exists a constant $C \in \mathbb{N}$ such that for all $\mathcal{F} \in \mathcal{P}$ and all $u, v \in \mathcal{F}$, $\mathrm{dist}_{\mathcal{F}(\mathcal{M})}(u, v) \leq C \cdot \|u - v\|$. The set $\mathcal{M}$ is $\|\cdot\|$-*norm-reducing* for $\mathcal{P}$ if for all $\mathcal{F} \in \mathcal{P}$ and all $u, v \in \mathcal{F}$ there exists $m \in \mathcal{M}$ such that $u + m \in \mathcal{F}$ and $\|u + m - v\| < \|u - v\|$.

The property of being norm-like does not depend on the norm, whereas being norm-reducing does. Norm-reducing sets are always norm-like, and norm-like sets are in turn always Markov bases, but the reverse of both statements is false in general (Example 3.4 and Example 3.5). For collections $\mathcal{P}_A$ however, every Markov basis is norm-like (Proposition 3.7).

Example 3.4. For any $n \in \mathbb{N}$, consider the normal set $\mathcal{F}_n := ([2] \times [n] \times \{0\}) \cup \{(2, n, 1)\}$ with the Markov basis $\{(0, 1, 0), (0, 0, 1), (-1, 0, -1)\}$. The distance between $(1, 1, 0)$ and $(2, 1, 0)$ in $\mathcal{F}_n(\mathcal{M})$ is $2n$ and thus $\mathcal{M}$ is not norm-like for $\{\mathcal{F}_n : n \in \mathbb{N}\}$ (see also Figure 2).

Example 3.5. Let $d \in \mathbb{N}$ and consider $A := (1, \ldots, 1) \in \mathbb{Z}^{1 \times d}$, then the set $\mathcal{M} := \{e_1 - e_i : 2 \leq i \leq d\}$ is a Markov basis for the collection $\mathcal{P}_A$. However, $\mathcal{M}$ is not $\|\cdot\|_p$-norm-reducing for any $d \geq 3$ and any $p \in [1, \infty]$. For instance, consider e_2 and e_3 in $\mathcal{F}_{A,1}(\mathcal{M})$. The only move from $\mathcal{M}$ that can be applied on e_2 is $e_1 - e_2$, but $\|(e_2 + e_1 - e_2) - e_3)\|_p = \|e_2 - e_3\|_p$. On the other hand, in the case we cannot find a move that decreases the 1-norm of two nodes $u, v \in \mathcal{F}_{A,b}$ by 1, we can find instead two moves $m_1, m_2 \in \mathcal{M}$ such that $u + m_1, u + m_1 + m_2 \in \mathcal{F}_{A,b}$ and $\|u + m_1 + m_2 - v\| = \|u - v\| - 2$. Thus, the graph-distance of any two elements u and v in $\mathcal{F}_{A,b}(\mathcal{M})$ is at most $\|u - v\|_1$ and hence $\mathcal{M}$ is norm-like for $\mathcal{P}_A$.

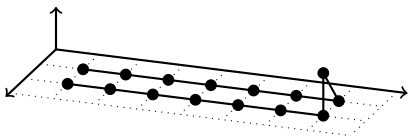

FIGURE 2. The graph from Example 3.4

Remark 3.6. Let $\mathcal{P}$ be a collection of finite subsets of $\mathbb{Z}^d$ and $\mathcal{M} \subset \mathbb{Z}^d$ be norm-like for $\mathcal{P}$. It follows from the definition that there exists a constant $C \in \mathbb{Q}_{\geq 0}$ such that for all $\mathcal{F} \in \mathcal{P}$

$$\mathrm{diam}(\mathcal{F}(\mathcal{M})) \leq C \cdot \max\{\|u - v\| : u, v \in \mathcal{F}\}.$$

The proof of our next results uses the *Graver basis* $\mathcal{G}_A \subset \mathbb{Z}^d$ for an integer matrix $A \in \mathbb{Z}^{m \times d}$ with $\ker_{\mathbb{Z}}(A) \cap \mathbb{N}^d = \{0\}$. We refer to [4, Chapter 3] for a precise definition.

Proposition 3.7. *Let $A \in \mathbb{Z}^{m \times d}$ with $\ker_{\mathbb{Z}}(A) \cap \mathbb{N}^d = \{0\}$ and $\mathcal{M} \subset \ker_{\mathbb{Z}}(A)$ be a Markov basis of $\mathcal{P}_A$. Then $\mathcal{M}$ is norm-like for $\mathcal{P}_A$.*

Proof. Let $\mathcal{M}$ be a Markov basis for $\mathcal{P}_A$. The Graver basis $\mathcal{G}_A$ for A is a finite set which is $\|\cdot\|_1$-norm-reducing for $\mathcal{P}_A$. Thus, define $C := \max_{g \in \mathcal{G}_A} \mathrm{diam}(\mathcal{F}_{A, Ag^+}(\mathcal{M}))$. Now, pick $u, v \in \mathcal{F}_{A,b}$ arbitrarily and let $u = v + \sum_{i=1}^{r} g_i$ be a walk from u to v in $\mathcal{F}_{A,b}(\mathcal{G}_A)$ of minimal length. Since the Graver basis is norm-reducing for $\mathcal{F}_{A,b}$, there always exists a path of length at most $\|u - v\|_1$ and hence $r \leq \|u - v\|_1$. Every g_i can be replaced by a path in $\mathcal{F}_{A, Ag_i^+}(\mathcal{M})$

of length at most C and these paths stay in $\mathcal{F}_{A,b}$. This gives a path of length $C \cdot r$, hence $\operatorname{dist}_{\mathcal{F}_{A,b}(\mathcal{M})}(u,v) \leq C\|u-v\|_1$. □

Proposition 3.8. *Let $\mathcal{P} \subset \mathbb{Z}^d$ be a polytope with $\dim(\mathcal{P} \cap \mathbb{Z}^d) > 0$ and let $\mathcal{M}$ be a Markov basis for $\mathcal{F}_i := (i \cdot \mathcal{P}) \cap \mathbb{Z}^d$ for all $i \in \mathbb{N}$. There exists a constant $C' \in \mathbb{Q}_{>0}$ such that for all $i \in \mathbb{N}$, $C' \cdot i \leq \operatorname{diam}(\mathcal{F}_i(\mathcal{M}))$. If $\mathcal{M}$ is norm-like for $\{\mathcal{F}_i : i \in \mathbb{N}\}$, then there exists a constant $C \in \mathbb{Q}_{>0}$ such that $\operatorname{diam}(\mathcal{F}_i(\mathcal{M})) \leq C \cdot i$ for all $i \in \mathbb{N}$.*

Proof. For the lower bound on the diameter, it suffices to show the existence of C' such that $C' \cdot i \leq \max\{\|u-v\| : u,v \in \mathcal{F}_i\}$ for all $i \in \mathbb{N}$ due to Lemma 3.1. Since $\dim(\mathcal{P} \cap \mathbb{Z}^d) > 0$, we can pick distinct $w, w' \in \mathcal{P} \cap \mathbb{Z}^d$. For all $i \in \mathbb{N}$, $i \cdot w, i \cdot w' \in \mathcal{F}_i$ and hence $i \cdot \|w - w'\| \leq \max\{\|u-v\| : u,v \in \mathcal{F}_i\}$.

To show the upper bound, assume that $\mathcal{M}$ is norm-like. It suffices to show that there exists $C \in \mathbb{Q}_{\geq 0}$ such that $\max\{\|u-v\| : u,v \in \mathcal{F}_i\} \leq i \cdot C$ by Remark 3.6. Now, let $v_1, \ldots, v_r \in \mathbb{Q}^d$ such that $\mathcal{P} = \operatorname{conv}_{\mathbb{Q}}(v_1, \ldots, v_r)$ and define $C := \max\{\|v_s - v_t\| : s \neq t\}$. Since $\mathcal{F}_i = (i \cdot \mathcal{P}) \cap \mathbb{Z}^d \subset \operatorname{conv}_{\mathbb{Q}}(iv_1, \ldots, iv_r)$ for all $i \in \mathbb{N}$, we have $\max\{\|u-v\| : u,v \in \mathcal{F}_i\} \leq \max\{\|iv_s - iv_t\| : s \neq t\} \leq C \cdot i$. □

Remark 3.9. Let $A \in \mathbb{Z}^{m \times n}$ with $\ker_{\mathbb{Z}}(A) \cap \mathbb{N}^d = \{0\}$ and let $\mathcal{M}$ be a Markov basis for $\mathcal{P}_A$. Then $\mathcal{M}$ is norm-like due to Proposition 3.7 and thus for all $b \in \mathbb{N}A$ there exists $C, C' \in \mathbb{Q}_{\geq 0}$ such that

$$i \cdot C' \leq \operatorname{diam}(\mathcal{F}_{A,ib}(\mathcal{M})) \leq i \cdot C$$

for all $i \in \mathbb{N}$. This generalizes for instance [20, Proposition 2.10] and [27, Example 4.7], where linear diameters on a ray in $\mathbb{N}A$ have been observed. This also implies that the construction of expanders from [27, Section 4] works for every right-hand side $b \in \mathbb{N}A$.

Remark 3.10. Let $A \in \mathbb{Z}^{m \times d}$ with $\ker_{\mathbb{Z}}(A) \cap \mathbb{N}^d = \{0\}$, $b \in \mathbb{N}A$, and let $\mathcal{M}$ be a Markov basis for $\mathcal{P}_A$. Proposition 3.8 provides a new proof that the simple walk on $(\mathcal{F}_{A,ib}(\mathcal{M}))_{i \in \mathbb{N}}$ cannot mix rapidly. The lower bound on the diameter from Proposition 3.8 implies, in general, the following upper bound on the edge-expansion (see for example [9, Proposition 1.30]):

$$h(\mathcal{F}_{A,i\cdot b}(\mathcal{M})) \leq |\mathcal{M}| \left(\exp\left(\frac{\log |\mathcal{F}_{A,i\cdot b}|}{D \cdot i} \right) - 1 \right).$$

In particular, the edge-expansion cannot be bounded from below by $\Omega(\frac{1}{p(i)})_{i \in \mathbb{N}}$ for a polynomial $p \in \mathbb{Q}[t]$ and since $(|\mathcal{F}_{A,i\cdot b}|)_{i \in \mathbb{N}} \in \mathcal{O}(i^r)_{i \in \mathbb{N}}$, the simple walk cannot mix rapidly. In [27], it was shown that the edge-expansion can be bounded from above by $\mathcal{O}(\frac{1}{i})_{i \in \mathbb{N}}$, which cannot be concluded from the upper expression.

We now turn our attention to the diameter of compressed fiber graphs. In particular, we want to know for which collections of normal sets is their diameter bounded. In general, compressing a fiber graph does not necessarily have an effect on the diameter (Example 3.11).

Although a low diameter is a necessary condition for good mixing, it is not sufficient. For instance, let G_n be the disjoint union of two complete graphs K_n connected by a single edge. Then $\operatorname{diam}(G_n) = 3$, but $h(G_n) \leq \frac{1}{n}$ implies that the simple walk does not mix rapidly.

Example 3.11. For any $n \in \mathbb{N}$, let $\mathcal{F}_n := \{(0,0),(0,1),(1,1),(1,2),\dots,(n,n)\} \subset \mathbb{Z}^2$. The unit vectors $\mathcal{M} = \{e_1, e_2\}$ are a Markov basis for $\{\mathcal{F}_n : n \in \mathbb{N}\}$. However, $\mathcal{F}_n^c(\mathcal{M}) = \mathcal{F}_n(\mathcal{M})$ and thus $\operatorname{diam}(\mathcal{F}_n^c(\mathcal{M})) = \operatorname{diam}(\mathcal{F}_n(\mathcal{M})) = 2n$ is unbounded.

Lemma 3.12. *Let $A \in \mathbb{Z}^{m\times d}$ and $z \in \ker_{\mathbb{Z}}(A)$. There exists $r \in [2d-2]$, distinct elements $g_1,\dots,g_r \in \mathcal{G}_A$, and $\lambda_1,\dots,\lambda_r \in \mathbb{N}_{>0}$ such that $z = \sum_{i=1}^r \lambda_i g_i$ and $g_i \sqsubseteq z$ for all $i \in [r]$*

Proof. This is [4, Lemma 3.2.3], although it only becomes clear from the original proof of [22, Theorem 2.1] that the appearing elements are all distinct. □

Proposition 3.13. *Let $A \in \mathbb{Z}^{m\times d}$ and $\mathcal{P} := \big\{\{x \in \mathbb{Z}^d : Ax = b, l \le x \le u\} : l,u \in \mathbb{Z}^d, b \in \mathbb{Z}^m\big\}$. Then for all $\mathcal{F} \in \mathcal{P}$, $\operatorname{diam}(\mathcal{F}^c(\mathcal{G}_A)) \le 2d-2$.*

Proof. Let $s,t \in \{x \in \mathbb{Z}^d : Ax = b, l \le x \le u\}$, then $s-t \in \ker_{\mathbb{Z}}(A)$ and thus $s = t + \sum_{i=1}^r \lambda_i g_i$ with $r \le 2d-2$, $\lambda_1,\dots,\lambda_r \in \mathbb{N}_{>0}$, and distinct $g_1,\dots,g_r \in \mathcal{G}_A$ such that $g_i \sqsubseteq s-t$ according to Lemma 3.12. It's now a consequence from [4, Lemma 3.2.4] that all intermediate points $t + \sum_{i=1}^k \lambda_i g_i$ for $k \le r$ are in $\{x \in \mathbb{Z}^d : Ax = b, l \le x \le u\}$. □

Lemma 3.14. *Let $\mathcal{F} \subset \mathbb{Z}^d$ be finite and let $\mathcal{F}_i := (i \cdot \operatorname{conv}_{\mathbb{Q}}(\mathcal{F})) \cap \mathbb{Z}^d$ for $i \in \mathbb{N}$. For all $u,v \in \mathcal{F}$, $\operatorname{dist}_{\mathcal{F}_i^c(\mathcal{M})}(iu,iv) \le \operatorname{dist}_{\mathcal{F}(\mathcal{M})}(u,v)$ for all $i \in \mathbb{N}$.*

Proof. The statement is trivially true if u and v are disconnected in $\mathcal{F}(\mathcal{M})$. Thus, assume the contrary and let $u = v + \sum_{j=1}^k m_j$ with $m_j \in \mathcal{M}$ be a path in $\mathcal{F}(\mathcal{M})$ of length $k = \operatorname{dist}_{\mathcal{F}(\mathcal{M})}(u,v)$ and let $i \in \mathbb{N}$. Clearly, $i \cdot u = i \cdot v + i \cdot \sum_{j=1}^k m_j = i \cdot v + \sum_{l=1}^k i \cdot m_j$, so it is left to prove that the elements traversed by this paths are in $\mathcal{F}_i$. Let $l \in [k]$, since $v + \sum_{j=1}^l m_j \in \mathcal{F}$, we have $i \cdot v + \sum_{j=1}^l i \cdot m_j \in i \cdot \mathcal{F} \subseteq \mathcal{F}_i$. Hence, this is a path in $\mathcal{F}_i^c(\mathcal{M})$ of length $k = \operatorname{dist}_{\mathcal{F}(\mathcal{M})}(u,v)$. □

We are ready to prove that the diameter of compressed fiber graphs coming from an integer matrix can be bounded for all right-hand sides simultaneously.

Theorem 3.15. *Let $A \in \mathbb{Z}^{m\times d}$ with $\ker_{\mathbb{Z}}(A) \cap \mathbb{N}^d = \{0\}$ and let $\mathcal{M}$ be a Markov basis for $\mathcal{P}_A$. There exists a constant $C \in \mathbb{N}$ such that $\operatorname{diam}(\mathcal{F}^c(\mathcal{M})) \le C$ for all $\mathcal{F} \in \mathcal{P}_A$.*

Proof. Our proof relies on basic properties of the Graver basis $\mathcal{G}_A$ of A. For any $g \in \mathcal{G}_A$, let $\mathcal{F}_g := \mathcal{F}_{A,Ag^+}$ and let $K := \max\{\operatorname{dist}_{\mathcal{F}_g(\mathcal{M})}(g^+,g^-) : g \in \mathcal{G}_A\}$. We show that the diameter of any compressed fiber graph of A is bounded from above by $(2d-2)\cdot K$. Let $b \in \mathbb{N}A$ arbitrary and choose elements $u,v \in \mathcal{F}_{A,b}$. According to Proposition 3.13, there exists $r \in [2d-2]$, $g_1,\dots,g_r \in \mathcal{G}_A$ and $\lambda_1,\dots,\lambda_r \in \mathbb{Z}$ such that $u = v + \sum_{i=1}^r \lambda_i g_i$, and $v + \sum_{i=1}^l \lambda_i g_i \in \mathbb{N}^d$ for all $l \in [r]$. According to Lemma 3.14, for any $i \in [r]$ there are $m_1^i,\dots,m_{k_i}^i \in \mathcal{M}$ and $\alpha_1,\dots,\alpha_{k_i} \in \mathbb{Z}$ such that $\lambda_i g_i^+ = \lambda_i g_i^- + \sum_{j=1}^{k_i} \alpha_j m_j^i$ is a path in the compression of $\mathcal{F}_{A,A\lambda_i g_i^+}(\mathcal{M})$ of length $k_i \le K$. Lifting these paths for every $i \in [r]$ yields a path $u = v + \sum_{i=1}^r \sum_{j=1}^{k_i} \alpha_j m_j^i$ in $\mathcal{F}_{A,b}^c(\mathcal{M})$ of length $r \cdot K \le (2d-2)\cdot K$. □

4. Heat-bath random walks

In this section, we establish the heat-bath random walk on compressed fiber graphs. We refer to [8] for a more general introduction on random walks of heat-bath type. Let $\mathcal{F} \subset \mathbb{Z}^d$ be finite set. For any $u \in \mathcal{F}$ and $m \in \mathbb{Z}^d$, the ray in $\mathcal{F}$ through u along m is denoted by $\mathcal{R}_{\mathcal{F},m}(u) := (u + m \cdot \mathbb{Z}) \cap \mathcal{F}$. Additionally, given a mass function $\pi : \mathcal{F} \to [0,1]$, we define

$$\mathcal{H}^{\pi}_{\mathcal{F},m}(x,y) := \begin{cases} \frac{\pi(y)}{\pi(\mathcal{R}_{\mathcal{F},m}(x))} & , \text{ if } y \in \mathcal{R}_{\mathcal{F},m}(x) \\ 0 & , \text{ otherwise} \end{cases}$$

for $x, y \in \mathcal{F}$. For $\mathcal{M} \subset \mathbb{Z}^d$ and a mass function $f : \mathcal{M} \to [0,1]$, the *heat-bath random walk* is

$$\mathcal{H}^{\pi,f}_{\mathcal{F},\mathcal{M}} = \sum_{m \in \mathcal{M}} f(m) \cdot \mathcal{H}^{\pi}_{\mathcal{F},m}. \tag{4.1}$$

The underlying graph of the heat-bath random walk is the compression $\mathcal{F}^c(\mathcal{M})$ and in this section, we assume throughout that for all $m \in \mathcal{M}$ and $\lambda \in \mathbb{Z} \setminus \{-1, 1\}$, $\lambda \cdot m \notin \mathcal{M}$. Let us first recall the basic properties of this random walk (compare also to [6, Lemma 2.2]).

Algorithm 1 Heat-bath random walk on compressed fiber graphs

Input: $\mathcal{F} \subset \mathbb{Z}^d$, $\mathcal{M} \subset \mathbb{Z}^d$, $v \in \mathcal{F}$, mass functions $f : \mathcal{M} \to [0,1]$ and $\pi : \mathcal{F} \to [0,1]$, $r \in \mathbb{N}$

1: **procedure** HeatBath:
2: $v_0 := v$
3: FOR $s = 0$; $s = s + 1$, $s < r$
4: Sample $m \in \mathcal{M}$ according to f
5: Sample $v_{s+1} \in \mathcal{R}_{\mathcal{F},m}(v_s)$ according to $\mathcal{R}_{\mathcal{F},m}(v_s) \to [0,1]$, $y \mapsto \frac{\pi(y)}{\pi(\mathcal{R}_{\mathcal{F},m}(v_s))}$
6: RETURN $v_1, \ldots, v_r$

Proposition 4.1. *Let $\mathcal{F} \subset \mathbb{Z}^d$ and $\mathcal{M} \subset \mathbb{Z}^d$ be finite sets. Let $f : \mathcal{M} \to [0,1]$ and $\pi : \mathcal{F} \to (0,1)$ be mass functions. Then $\mathcal{H}^{\pi,f}_{\mathcal{F},\mathcal{M}}$ is aperiodic, has stationary distribution π, is reversible with respect to π, and all of its eigenvalues are non-negative. The random walk is irreducible if and only if $\{m \in \mathcal{M} : f(m) > 0\}$ is a Markov basis for $\mathcal{F}$.*

Proof. Since for any $u \in \mathcal{F}$ and any $m \in \mathcal{M}$, $\mathcal{H}^{\pi}_{\mathcal{F},m}(u,u) > 0$, there are halting states and thus $\mathcal{H}^{\pi,f}_{\mathcal{F},\mathcal{M}}$ is aperiodic. By definition, $\pi(x)\mathcal{H}^{\pi}_{\mathcal{F},m}(x,y) = \pi(y)\mathcal{H}^{\pi}_{\mathcal{F},m}(y,x)$ and thus $\mathcal{H}^{\pi,f}_{\mathcal{F},\mathcal{M}}$ is reversible with respect to π and π is a stationary distribution. The statement on the eigenvalues is exactly [8, Lemma 1.2]. Let $\mathcal{M}' = \{m \in \mathcal{M} : f(m) > 0\}$ and $f' = f|_{\mathcal{M}'}$, then $\mathcal{H}^{\pi,f}_{\mathcal{F},\mathcal{M}} = \mathcal{H}^{\pi,f'}_{\mathcal{F},\mathcal{M}'}$ and thus the heat-bath random walk is irreducible if and only if $\mathcal{M}'$ is a Markov basis for $\mathcal{F}$. □

Remark 4.2. Analyzing the speed of convergence of random walks with second largest eigenvalues does not take the computation time of a single transition into account. From a computational point of view, the difference of the simple walk and the heat-bath random walk is Step 4 of Algorithm 1. However, we argue that Step 4 can be done efficiently in many cases. For instance, a hard normalizing constant of π cancels out. If π is the uniform distribution, then one needs to sample uniformly from $\mathcal{R}_{\mathcal{F},m}(v)$ in Step 4, which can be done

efficiently. If the input of Algorithm 1 is a normal set $\mathcal{F} = \{u \in \mathbb{Z}^d : Au \leq b\}$ that is given in $\mathcal{H}$-representation, then the length of the ray $\mathcal{R}_{\mathcal{F},m}(v)$ can be computed with a number of rounding, division, and comparing operations that is linear in the number of rows of A.

There are situations in which the heat-bath random walk provides no speed-up compared with the simple walk (Example 4.3). Intuitively, adding more moves to the set of allowed moves should improve the mixing time of the random walk. In general, however, this is not true for the heat-bath walk (Example 4.4).

Example 4.3. For $n \in \mathbb{N}$, consider the normal set

$$\mathcal{F}_n := \left\{ \begin{bmatrix} 0 & 1 & 1 & \cdots & 1 \\ 1 & 0 & 0 & \cdots & 0 \end{bmatrix}, \begin{bmatrix} 1 & 0 & 1 & \cdots & 1 \\ 0 & 1 & 0 & \cdots & 0 \end{bmatrix}, \ldots, \begin{bmatrix} 1 & 1 & \cdots & 1 & 0 \\ 0 & 0 & \cdots & 0 & 1 \end{bmatrix} \right\} \subset \mathbb{Q}^{2\times n}.$$

In the language of [7, Section 1.1], $\mathcal{F}_n$ is precisely the fiber of the $2 \times n$ independence model where row sums are $(n-1, 1)$ and column sums are $(1, 1, \ldots, 1)$. The minimal Markov basis of the independence model, often referred to as the *basic moves*, is precisely the set $\mathcal{M}_n := \{v - u : u, v \in \mathcal{F}_n\} \setminus \{0\}$. In particular, the fiber graph $\mathcal{F}_n(\mathcal{M}_n)$ is the complete graph on n nodes. All rays along basic moves have length 2 and thus the transition matrices of the simple random walk and the heat-bath random walk coincide. There are $n \cdot (n-1)$ many basic moves and the transition matrix of both random walks is

$$\frac{1}{n(n-1)} \begin{bmatrix} 1 & \ldots & 1 \\ \vdots & & \vdots \\ 1 & \ldots & 1 \end{bmatrix} + \frac{(n(n-1)-n)}{n(n-1)} \cdot I_n.$$

The second largest eigenvalue is $1 - \frac{1}{n-1}$ which implies that for $n \to \infty$, neither the simple walk nor the heat-bath random walk are rapidly mixing.

Example 4.4. Let $\mathcal{F} = [2] \times [5] \subset \mathbb{Z}^2$, $\mathcal{M} = \{e_1, e_2, 2e_1 + e_2\}$, and let π be the uniform distribution on $\mathcal{F}$. Since $\{e_2, 2e_1 + e_2\}$ is not a Markov basis for $\mathcal{F}$, any mass function $f : \mathcal{M} \to [0, 1]$ must have $f(e_1) > 0$ in order to make the corresponding heat-bath random walk irreducible. Comparing the second largest eigenvalue modulus of heat-bath random walks that sample from $\{e_1, e_2\}$ and $\mathcal{M}$ uniformly, we obtain

$$\lambda\left(\frac{1}{2}\mathcal{H}^{\pi}_{\mathcal{F},e_1} + \frac{1}{2}\mathcal{H}^{\pi}_{\mathcal{F},e_2}\right) = \frac{1}{2} < \frac{2}{3} = \lambda\left(\frac{1}{3}\mathcal{H}^{\pi}_{\mathcal{F},e_1} + \frac{1}{3}\mathcal{H}^{\pi}_{\mathcal{F},e_2} + \frac{1}{3}\mathcal{H}^{\pi}_{\mathcal{F},2e_1+e_2}\right).$$

So, adding $2e_1 + e_2$ to the set of allowed moves slows the walk down. This phenomenon does not appear for the simple walk on $\mathcal{F}$, where the second largest eigenvalue modulus improves from ≈ 0.905 to ≈ 0.888 when adding the move $2e_1 + e_2$ to the Markov basis.

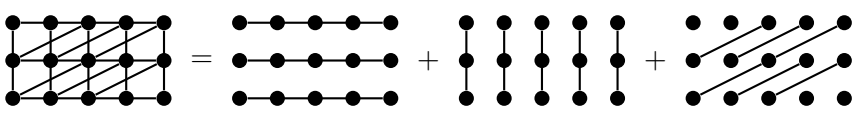

FIGURE 3. Decomposition of the graph in Example 4.4

Remark 4.5. Let $\mathcal{F} \subset \mathbb{Z}^d$ be finite and $\mathcal{M} = \{m_1, \dots, m_d\} \subset \mathbb{Z}^d$ be a linearly independent Markov basis of $\mathcal{F}$. If the moves are selected uniformly, then the heat-bath random walk on $\mathcal{F}$ coincides with the *Glauber dynamics* on $\mathcal{F}$. To see it, choose $u \in \mathcal{F}$ and let

$$\mathcal{F}' := \{\lambda \in \mathbb{Z}^d : u + \lambda_1 m_1 + \cdots + \lambda_d m_d \in \mathcal{F}\}.$$

It is easy to check that $\mathcal{F}'$ is unique up to translation and depends only on $\mathcal{F}$, u, and $\mathcal{M}$. Since the vectors in $\mathcal{M}$ are linearly independent, every element of $\mathcal{F}$ can be represented by a unique choice of coefficients in $\mathcal{F}'$. Thus, the heat-bath random walk on $\mathcal{F}$ using $\mathcal{M}$ is equivalent to the heat-bath random walk on on $\mathcal{F}'$ using the unit vectors as moves. For any unit vector $e_i \in \mathbb{Z}^d$, the ray through an element $v \in \mathcal{F}'$ is $\{w \in \mathcal{F} : w_j = v_j \forall j \neq i\}$ and this is precisely the form desired in the Glauber dynamics [14, Section 3.3.2].

For the remainder of this section, we primarily focus on heat-bath random walks $\mathcal{H}^{\pi,f}_{\mathcal{F},\mathcal{M}}$ that converge to the uniform distribution π on a finite, but not necessarily normal, set $\mathcal{F}$. We particularly aim for bounds on its second largest eigenvalue by making use of the decomposition from equation 4.1. Our first observations consider its summands $\mathcal{H}^{\pi}_{\mathcal{F},m}$ that can be well understood analytically (Proposition 4.6) and combinatorially (Proposition 4.7).

Proposition 4.6. *Let $\mathcal{F} \subset \mathbb{Z}^d$ be a finite set, $m \in \mathbb{Z}^d$, and $\pi : \mathcal{F} \to [0,1]$ be the uniform distribution. Let $\mathcal{R}_1, \dots, \mathcal{R}_k$ be the disjoint rays through $\mathcal{F}$ along m. Then*

1. *$\mathcal{H}^{\pi}_{\mathcal{F},m}$ is symmetric and idempotent.*
2. *$\operatorname{img}(\mathcal{H}^{\pi}_{\mathcal{F},m}) = \operatorname{span}_{\mathbb{R}} \left\{ \sum_{x \in \mathcal{R}_1} e_x, \sum_{x \in \mathcal{R}_2} e_x, \dots, \sum_{x \in \mathcal{R}_k} e_x \right\}$.*
3. *$\ker(\mathcal{H}^{\pi}_{\mathcal{F},m}) = \bigoplus_{i=1}^{k} \operatorname{span}_{\mathbb{R}} \{e_x - e_y : x, y \in \mathcal{R}_i, x \neq y\}$.*
4. *$\operatorname{rank}(\mathcal{H}^{\pi}_{\mathcal{F},m}) = k$ and $\dim \ker(\mathcal{H}^{\pi}_{\mathcal{F},m}) = |\mathcal{F}| - k$.*
5. *The spectrum of $\mathcal{H}^{\pi}_{\mathcal{F},m}$ is $\{0,1\}$.*

Proof. Symmetry of $\mathcal{H}^{\pi}_{\mathcal{F},m}$ follows from the definition. By assumption, $\mathcal{F}$ is the disjoint union of $\mathcal{R}_1, \dots, \mathcal{R}_k$ and hence there exists a permutation matrix S such that $S\mathcal{H}^{\pi}_{\mathcal{F},m}S^T$ is a block matrix whose building blocks are the matrices

$$\frac{1}{|\mathcal{R}_i|} \begin{bmatrix} 1 & \dots & 1 \\ \vdots & & \vdots \\ 1 & \dots & 1 \end{bmatrix} \in \mathbb{Q}^{|\mathcal{R}_i| \times |\mathcal{R}_i|}.$$

Thus, $\mathcal{H}^{\pi}_{\mathcal{F},m}$ is idempotent and the rank of $\mathcal{H}^{\pi}_{\mathcal{F},m}$ is k. A basis of its image and its kernel can be read off directly and idempotent matrices can only have the eigenvalues 0 and 1. □

Proposition 4.7. *Let $\mathcal{F} \subset \mathbb{Z}^d$ and $\mathcal{M} \subset \mathbb{Z}^d$ be finite sets, $\pi : \mathcal{F} \to [0,1]$ be the uniform distribution, and let $V_1, \dots, V_c \subseteq \mathcal{F}$ be the nodes of the connected components of $\mathcal{F}(\mathcal{M})$, then*

$$\bigcap_{m \in \mathcal{M}} \operatorname{img}(\mathcal{H}^{\pi}_{\mathcal{F},m}) = \operatorname{span}_{\mathbb{R}} \left\{ \sum_{x \in V_1} e_x, \dots, \sum_{x \in V_c} e_x \right\}.$$

Proof. It is clear by Proposition 4.6 that the set on the right-hand side is contained in any $\operatorname{img}(\mathcal{H}^{\pi}_{\mathcal{F},m})$ since any V_i decomposes disjointly into rays along $m \in \mathcal{M}$. To show the other inclusion, write $\mathcal{M} = \{m_1, \dots, m_k\}$ and let for any $i \in [k]$, $\mathcal{R}^i_1, \dots, \mathcal{R}^i_{n_i}$ be the disjoint rays

through $\mathcal{F}$ parallel to m_i. In particular, $\{\mathcal{R}^i_1, \dots, \mathcal{R}^i_{n_i}\}$ is a partition of $\mathcal{F}$ for any $i \in [k]$. Let $v \in \bigcap_{m\in\mathcal{M}} \operatorname{img}(\mathcal{H}^\pi_{\mathcal{F},m})$. Again by Proposition 4.6, there exists for any $i \in [k]$, $\lambda^i_1, \dots, \lambda^i_{n_i} \in \mathbb{Q}$ such that

$$v = \sum_{j=1}^{n_i} \sum_{x\in\mathcal{R}^i_j} \lambda^i_j e_x.$$

Notice that if two distinct Markov moves m_i and $m_{i'}$ and two indices $j \in [n_i]$ and $j' \in [n_{i'}]$ satisfy $\mathcal{R}^i_j \cap \mathcal{R}^{i'}_{j'} \neq \emptyset$, then $\lambda^i_j = \lambda^{i'}_{j'}$. We show that for any $i \in [k]$ and any $a \in [c]$, $\lambda^i_j = \lambda^i_{j'}$ when $\mathcal{R}^i_j$ and $\mathcal{R}^i_{j'}$ are a subset of V_a. This implies the proposition. So take distinct $x, x' \in V_a$ and assume that x and x' lie on different rays of m_i and let that be $x \in \mathcal{R}^i_j$ and $x' \in \mathcal{R}^i_{j'}$ with $j \neq j'$. Since x and x' are in the same connected component V_a of $\mathcal{F}(\mathcal{M})$, let $y_{i_0}, \dots, y_{i_r} \in \mathcal{F}$ be the nodes on a minimal path in $\mathcal{F}^c(\mathcal{M})$ with $y_{i_0} = x$ and $y_{i_r} = x'$. For any $s \in [r]$, y_{i_s} and $y_{i_{s-1}}$ are contained in the same ray $\mathcal{R}^{k_s}_{t_s}$ coming from a Markov move m_{k_s}. In particular, $\mathcal{R}^{t_{s-1}}_{k_{s-1}} \cap \mathcal{R}^{k_s}_{t_s} \neq \emptyset$ and due to our observation made above $\lambda^i_j = \lambda^{k_1}_{t_1} = \lambda^{k_2}_{t_2} = \cdots = \lambda^{k_r}_{t_r} = \lambda^i_{j'}$ which finishes the proof. □

Definition 4.8. Let $\mathcal{F} \subset \mathbb{Z}^d$ and $\mathcal{M} \subset \mathbb{Z}^d$ be finite sets and $\mathcal{M}' \subseteq \mathcal{M}$. Let $\mathcal{V}$ be the set of connected components of $\mathcal{F}(\mathcal{M}\setminus\mathcal{M}')$ and $\mathcal{R}$ be the set of all rays through $\mathcal{F}$ along all elements of $\mathcal{M}'$. The *ray matrix* of $\mathcal{F}(\mathcal{M})$ along $\mathcal{M}'$ is $A_\mathcal{F}(\mathcal{M}, \mathcal{M}') := (|R \cap V|)_{R\in\mathcal{R}, V\in\mathcal{V}} \in \mathbb{N}^{\mathcal{R}\times\mathcal{V}}$.

Example 4.9. Let $\mathcal{F} = [3] \times [3]$, $\mathcal{M} = \{e_1, e_2, e_1 + e_2\}$, and $\mathcal{M}' = \{e_1, e_2\}$. Then $\mathcal{F}(\mathcal{M} \setminus \mathcal{M}')$ has five connected components and the ray matrix of $\mathcal{F}(\mathcal{M})$ along $\mathcal{M}'$ is

$$A_\mathcal{F}(\mathcal{M}, \mathcal{M}') = \begin{bmatrix} 1 & 1 & 1 & 0 & 0 \\ 0 & 1 & 1 & 1 & 0 \\ 0 & 0 & 1 & 1 & 1 \\ 0 & 0 & 1 & 1 & 1 \\ 0 & 1 & 1 & 1 & 0 \\ 1 & 1 & 1 & 0 & 0 \end{bmatrix}.$$

Remark 4.10. Let $\mathcal{F} \subset \mathbb{Z}^2$, then the rays through $\mathcal{F}$ along e_1 are the connected components of $\mathcal{F}(\{e_1, e_2\} \setminus \{e_2\})$ and the rays through $\mathcal{F}$ along e_2 are the connected components of $\mathcal{F}(\{e_1, e_2\} \setminus \{e_1\})$, thus $A_\mathcal{F}(\mathcal{M}, e_1) = A_\mathcal{F}(\mathcal{M}, e_2)^T$.

Proposition 4.11. *Let $\mathcal{F} \subset \mathbb{Z}^d$ and $\mathcal{M} \subset \mathbb{Z}^d$ be finite sets, $\pi : \mathcal{F} \to [0, 1]$ be the uniform distribution, and $\mathcal{M}' \subseteq \mathcal{M}$. Then*

$$\ker(A_\mathcal{F}(\mathcal{M}, \mathcal{M}')) \cong \bigcap_{m\in\mathcal{M}\setminus\mathcal{M}'} \operatorname{img}(\mathcal{H}^\pi_{\mathcal{F},m}) \cap \bigcap_{m\in\mathcal{M}'} \ker(\mathcal{H}^\pi_{\mathcal{F},m}).$$

Proof. Let $V_1, \dots, V_c$ be the connected components of $\mathcal{F}(\mathcal{M} \setminus \mathcal{M}')$ and $\mathcal{R}_1, \dots, \mathcal{R}_r$ be the rays along elements in $\mathcal{M}'$. Let $I := \bigcap_{m\in\mathcal{M}\setminus\mathcal{M}'} \operatorname{img}(\mathcal{H}^\pi_{\mathcal{F},m})$ and $K := \bigcap_{m\in\mathcal{M}'} \ker(\mathcal{H}^\pi_{\mathcal{F},m})$. By Proposition 4.7, any element of I has the form $v = \sum_{i=1}^c (\lambda_i \sum_{x\in V_i} e_x)$ for $\lambda_1, \dots, \lambda_c \in \mathbb{Q}$. Assume additionally that $v \in \ker(\mathcal{H}^\pi_{\mathcal{F},m})$ for $m \in \mathcal{M}'$ and let $\mathcal{R}_{i_1}, \dots \mathcal{R}_{i_j}$ be the rays which belong to m, then for any $k \in [j]$, $0 = \sum_{x\in\mathcal{R}_{i_k}} v_x = \sum_{j=1}^c \lambda_j |\mathcal{R}_{i_k} \cap V_j|$. Put differently, a vector $\lambda \in \mathbb{R}^c$ is in the kernel of $(|\mathcal{R}_i \cap V_j|)_{i\in[r], j\in[c]}$ if and only if $\sum_{i=1}^c (\lambda_i \sum_{x\in V_i} e_x) \in I \cap K$. □

Conditions on the kernel of the ray matrix allow us to give a lower bound on the second largest eigenvalue of the heat-bath random walk.

Proposition 4.12. *Let $\mathcal{F} \subset \mathbb{Z}^d$ and $\mathcal{M} \subset \mathbb{Z}^d$ be finite sets and π be the uniform distribution. Let $\mathcal{M}' \subseteq \mathcal{M}$ such that $\ker(A_{\mathcal{F}}(\mathcal{M}, \mathcal{M}')) \neq \{0\}$, then $\lambda(\mathcal{H}^{\pi,f}_{\mathcal{F},\mathcal{M}}) \geq 1 - \sum_{m\in\mathcal{M}'} f(m)$ for any mass function $f : \mathcal{M} \to [0,1]$.*

Proof. Using the isomorphism from Proposition 4.11, we can choose a non-zero $v \in \mathbb{Q}^P$ such that $\mathcal{H}^{\pi}_{\mathcal{F},m} v = v$ for all $m \in \mathcal{M} \setminus \mathcal{M}'$ and $\mathcal{H}^{\pi}_{\mathcal{F},m} v = 0$ for all $m \in \mathcal{M}'$. In particular

$$\mathcal{H}^{\pi,f}_{\mathcal{F},\mathcal{M}} v = \sum_{m\in\mathcal{M}} f(m)\mathcal{H}^{\pi}_{\mathcal{F},m} v = \sum_{m\in\mathcal{M}\setminus\mathcal{M}'} f(m)\mathcal{H}^{\pi}_{\mathcal{F},m} v = \sum_{m\in\mathcal{M}\setminus\mathcal{M}} f(m) v.$$

Since f is a mass function, $1 - \sum_{m\in\mathcal{M}'} f(m)$ is an eigenvalue of $\mathcal{H}^{\pi,f}_{\mathcal{F},\mathcal{M}}$. □

Definition 4.13. Let $\mathcal{F} \subset \mathbb{Z}^d$ and $m, m' \in \mathbb{Z}^d$ not collinear. The pair (m, m') has the *intersecting ray property* in $\mathcal{F}$ if the following holds: For any pair of rays $\mathcal{R}_1, \mathcal{R}_2$ parallel to m and any pair of rays $\mathcal{R}'_1, \mathcal{R}'_2$ parallel to m' where both $\mathcal{R}_1 \cap \mathcal{R}'_1$ and $\mathcal{R}_2 \cap \mathcal{R}'_2$ are not empty, then $\mathcal{R}_1 \cap \mathcal{R}'_2 \neq \emptyset$ implies $\mathcal{R}'_1 \cap \mathcal{R}_2 \neq \emptyset$ and $|\mathcal{R}_1| \cdot |\mathcal{R}'_1|^{-1} = |\mathcal{R}_2| \cdot |\mathcal{R}'_2|^{-1}$. For a finite set $\mathcal{M} \subset \mathbb{Z}^d$, the graph $\mathcal{F}^c(\mathcal{M})$ has the *intersecting ray property* if all (m, m') have the intersecting ray property in $\mathcal{F}$.

Example 4.14. The compressed fiber graph on $[n_1] \times \cdots \times [n_d] \subset \mathbb{Z}^d$ that uses the unit vectors $\{e_1, \ldots, e_d\}$ as moves has the intersecting ray property. On the other hand, consider $\mathcal{F} = \{u \in \mathbb{N}^2 : u_1 + u_2 \leq 1\}$ and take the rays $\mathcal{R}_1 := \{(0,0), (0,1)\}$ and $\mathcal{R}_2 := \{(1,0)\}$ that are parallel to e_2 and the rays $\mathcal{R}'_1 := \{(0,1)\}$ and $\mathcal{R}'_2 := \{(0,0), (1,0)\}$ that are parallel to e_1. Then $\mathcal{R}_1 \cap \mathcal{R}'_1 = \{(1,0)\}$ and $\mathcal{R}_2 \cap \mathcal{R}'_2 = \{(0,1)\}$, but $\mathcal{R}_1 \cap \mathcal{R}'_2 = \{(0,0)\} \neq \emptyset$ and $\mathcal{R}'_1 \cap \mathcal{R}_2 = \emptyset$.

Proposition 4.15. *Let $m, m' \in \mathbb{Z}^d$ not collinear and $\mathcal{F} \subset \mathbb{Z}^d$ be a finite set. The matrices $\mathcal{H}^{\pi}_{\mathcal{F},m}$ and $\mathcal{H}^{\pi}_{\mathcal{F},m'}$ commute if and only if (m, m') have the intersecting ray property in $\mathcal{F}$.*

Proof. Let $u_1, u_2 \in \mathcal{F}$. Then

$$(\mathcal{H}^{\pi}_{\mathcal{F},m} \cdot \mathcal{H}^{\pi}_{\mathcal{F},m'})_{u_1,u_2} = \begin{cases} |\mathcal{R}_{\mathcal{F},m}(u_1)|^{-1} \cdot |\mathcal{R}_{\mathcal{F},m'}(u_2)|^{-1}, & \text{if } \mathcal{R}_{\mathcal{F},m}(u_1) \cap \mathcal{R}_{\mathcal{F},m'}(u_2) \neq \emptyset \\ 0, & \text{otherwise} \end{cases}.$$

Let $\mathcal{R}_1 := \mathcal{R}_{\mathcal{F},m}(u_1)$, $\mathcal{R}'_1 := \mathcal{R}_{\mathcal{F},m'}(u_1)$, $\mathcal{R}_2 := \mathcal{R}_{\mathcal{F},m}(u_2)$, and $\mathcal{R}'_2 := \mathcal{R}_{\mathcal{F},m'}(u_2)$ Thus, $(\mathcal{H}^{\pi}_{\mathcal{F},m} \cdot \mathcal{H}^{\pi}_{\mathcal{F},m'})_{u_1,u_2} = (\mathcal{H}^{\pi}_{\mathcal{F},m'} \cdot \mathcal{H}^{\pi}_{\mathcal{F},m})_{u_1,u_2}$. It is easy to see that the matrices commute if and only if (m, m') have the intersecting ray property. □

Lemma 4.16. *Let $H_1, \ldots, H_n \in \mathbb{R}^{n\times n}$ be pairwise commuting matrices. Then any eigenvalue of $\sum_{i=1}^n H_i$ has the form $\lambda_1 + \cdots + \lambda_n$ where λ_i is an eigenvalue of H_i.*

Proof. This is a straightforward extension of the case $n = 2$ in [12, Theorem 2.4.8.1] and relies on the fact that commuting matrices are simultaneously triangularizable. □

Proposition 4.17. *Let $\mathcal{F} \subset \mathbb{Z}^d$ and $\mathcal{M} \subset \mathbb{Z}^d$ be finite sets and suppose there exists $m \in \mathcal{M}$ such that (m, m') has the intersecting ray property in $\mathcal{F}$ for all $m' \in \mathcal{M}' := \mathcal{M} \setminus \{m\}$. Let $\mathcal{V}_1, \dots, \mathcal{V}_c$ be the connected components of $\mathcal{F}(\mathcal{M}')$, $\pi_i : \mathcal{V}_i \to [0,1]$ the uniform distribution, and $f' = (1 - f(m))^{-1} \cdot f|_{\mathcal{M}'}$, then*

$$\lambda(\mathcal{H}^{\pi,f}_{\mathcal{F},\mathcal{M}}) \leq f(m) + (1 - f(m)) \cdot \max\{\lambda(\mathcal{H}^{\pi_i,f'}_{\mathcal{V}_i,\mathcal{M}'}) : i \in [c]\}.$$

Proof. Let $\mathcal{H} := \mathcal{H}^{\pi,f'}_{\mathcal{F},\mathcal{M}'}$ be the heat-bath random walk on $\mathcal{F}(\mathcal{M})$ that samples moves from $\mathcal{M}'$ according to f', then $\mathcal{H}^{\pi,f}_{\mathcal{F},\mathcal{M}} = f(m) \cdot \mathcal{H}^{\pi}_{\mathcal{F},m} + (1 - f(m)) \cdot \mathcal{H}$. By assumption, all pairs (m, m') with $m' \in \mathcal{M}'$ have the intersecting ray property and thus the matrices $\mathcal{H}^{\pi}_{\mathcal{F},m}$ and $\mathcal{H}$ commute according to Proposition 4.15. The eigenvalues of all involved matrices are non-negative and thus Lemma 4.16 implies that the second largest eigenvalue of $\mathcal{H}^{\pi,f}_{\mathcal{F},\mathcal{M}}$ has the form $\lambda + \lambda'$ where $\lambda \in \{0, f(m)\}$ by Proposition 4.6 and where λ' is an eigenvalue of $(1 - f(m)) \cdot \mathcal{H}$. The matrix $\mathcal{H}$ is a block matrix whose building blocks are the matrices $\mathcal{H}^{\pi,f'}_{\mathcal{V}_i,\mathcal{M}'} = \mathcal{H}^{\pi_i,f'}_{\mathcal{V}_i,\mathcal{M}'}$ and thus the statement follows. □

Proposition 4.18. *Let $\mathcal{F} \subset \mathbb{Z}^d$ and $\mathcal{M} \subset \mathbb{Z}^k$ be finite sets. If $\mathcal{F}(\mathcal{M})$ has the intersecting ray property, then $\lambda(\mathcal{H}^{\pi,f}_{\mathcal{F},\mathcal{M}}) \leq 1 - \min(f)$.*

Proof. Let $\mathcal{M} = \{m_1, \dots, m_k\}$. The intersecting ray property and Proposition 4.15 give that the matrices $f(m_1) \cdot \mathcal{H}^{\pi}_{\mathcal{F},m_i}, \dots, f(m_k) \cdot \mathcal{H}^{\pi}_{\mathcal{F},m_k}$ commute pairwise. According to Proposition 4.6, the eigenvalues of $f(m_i) \cdot \mathcal{H}^{\pi}_{\mathcal{F},m_i}$ are $\{0, f(m_i)\}$. Lemma 4.16 gives that the second largest eigenvalue of $\mathcal{H}^{\pi,f}_{\mathcal{F},\mathcal{M}}$, which equals the second largest eigenvalue modulus since all of its eigenvalues are non-negative, fulfills $\lambda(\mathcal{H}^{\pi,f}_{\mathcal{F},\mathcal{M}}) = \sum_{i \in I} f(m_i)$ for a subset $I \subseteq [k]$. Since $\lambda(\mathcal{H}^{\pi,f}_{\mathcal{F},\mathcal{M}}) < 1$ and $\sum_{i=1}^{k} f(m_i) = 1$, we have $I \neq [k]$ and the claim follows. □

Proposition 4.19. *Let $n_1, \dots, n_d \in \mathbb{N}_{>1}$, $\mathcal{F} = [n_1] \times \cdots \times [n_d]$, and $\mathcal{M} = \{e_1, \dots, e_d\}$. Then for any positive mass function $f : \mathcal{M} \to [0,1]$, $\lambda(\mathcal{H}^{\pi,f}_{\mathcal{F},\mathcal{M}}) = 1 - \min(f)$.*

Proof. It is easy to verify that $\mathcal{F}^c(\mathcal{M})$ has the intersecting ray property and thus Proposition 4.18 shows $\lambda(\mathcal{H}^{\pi,f}_{\mathcal{F},\mathcal{M}}) \leq 1 - \min(f)$. Assume that $\min(f) = f(e_i)$. The connected components of $\mathcal{F}^c(\{e_1, \dots, e_d\} \setminus \{e_i\})$ are the layers $V_j := \{u \in \mathcal{F} : u_i = j\}$ for any $j \in [n_i]$ and the rays through $\mathcal{F}$ parallel are $\mathcal{R}_k := \{(0,k) + s \cdot e_i : s \in [n_i]\}$ for $k = (k_1, \dots, k_{i-1}, k_{i+1}, \dots, k_d) \in [n_1] \times \cdots \times [n_{i-1}] \times [n_{i+1}] \times \cdots \times [n_d]$. In particular, any ray intersects any connected component exactly once. Thus, the matrix $(|R_k \cap V_j|)_{k,j}$ is the all-ones matrix, which has a non-trivial kernel. Proposition 4.12 implies $\lambda(\mathcal{H}^{\pi,f}_{\mathcal{F},\mathcal{M}}) \geq 1 - f(e_i)$. □

Remark 4.20. In the special case $n := n_1 = \cdots = n_d$ and $f : \{e_1, \dots, e_d\} \to [0,1]$ the uniform distribution in Proposition 4.19, the heat-bath random walk on $[n]^d$ is known as *Rook's walk* in the literature. In this case, Proposition 4.19 is exactly [13, Proposition 2.3]. In [18], upper bounds on the mixing time of the Rook's walk were obtained with *path-coupling*.

The stationary distribution of the heat-bath random walk is independent of the actual mass function on the Markov moves. The problem of finding the mass function which leads

to the fastest mixing behaviour can be formulated as the following optimization problem:

$$\arg\min\left\{\lambda(\mathcal{H}^{\pi,f}_{\mathcal{F},\mathcal{M}}) : f : \mathcal{M} \to (0,1),\ \sum_{m\in\mathcal{M}} f(m) = 1\right\}. \tag{4.2}$$

It follows from Proposition 4.19 that the optimal value of (4.2) for $\mathcal{F} = [n_1] \times \cdots \times [n_d]$, $\mathcal{M} = \{e_1, \ldots, e_d\}$, and the uniform distribution π on $\mathcal{F}$ is the uniform distribution on $\mathcal{M}$. Another example where the uniform distribution is the optimal solution to (4.2), but where the verification is more involved, is presented in Example 4.21.

Example 4.21. Let $\mathcal{F} = [2] \times [5]$ as in Example 4.4 and consider $\mathcal{M} = \{e_1, 2e_1 + e_2\}$. We investigate for which $\mu \in (0,1)$, the transition matrix $\mu\mathcal{H}^{\pi}_{\mathcal{F},e_1} + (1-\mu)\mathcal{H}^{\pi}_{\mathcal{F},2e_1+e_2}$ has the smallest second largest eigenvalue modulus. Its characteristic polynomial in $\mathbb{Q}[\mu, x]$ is

$$-\frac{1}{25}x^4(x-1)(\mu+x-1)^6(-5x^2+5x+2\mu^2-2\mu)(-5x^2+5x+4\mu^2-4\mu)$$

and hence its eigenvalues are

$$x_1(\mu) := 1, \quad x_2(\mu) := 1-\mu,$$
$$x_3(\mu) := \frac{1}{2}\left[1+\sqrt{1+\frac{8}{5}(\mu^2-\mu)}\right], \quad x_4(\mu) := \frac{1}{2}\left[1-\sqrt{1+\frac{8}{5}(\mu^2-\mu)}\right],$$
$$x_5(\mu) := \frac{1}{2}\left[1+\sqrt{1+4(\mu^2-\mu)}\right], \quad x_6(\mu) := \frac{1}{2}\left[1-\sqrt{1+4(\mu^2-\mu)}\right].$$

It is straightforward to check that $x_5(\mu) > \frac{1}{2} > x_6(\mu)$, $x_3(\mu) > \frac{1}{2} > x_4(\mu)$. Since $\mu^2 - \mu < 0$ for $u \in (0,1)$ and $x_3(\mu) \geq x_6(\mu)$. We can show that $x_4(\mu) \geq x_2(\mu)$ and thus

$$\lambda(\mu\mathcal{H}^{\pi}_{\mathcal{F},e_1} + (1-\mu)\mathcal{H}^{\pi}_{\mathcal{F},2e_1+e_2}) = \frac{1}{2}\left[1+\sqrt{1+\frac{8}{5}(\mu^2-\mu)}\right].$$

The fastest heat-bath random walk on $\mathcal{F}(\mathcal{M})$ which converges to uniform is thus obtained for $\mu = \frac{1}{2}$, i.e. when the moves are selected uniformly. The second largest eigenvalue in this case is $\frac{1}{10}(5+\sqrt{15}) \approx 0.887$, which is larger than the second largest eigenvalue of the heat-bath walk that selects uniformly from $\{e_1, e_2\}$ (see Proposition 4.19).

5. Augmenting Markov bases

It follows from our investigation in Section 3 that the diameter of all compressed fiber graphs coming from a fixed integer matrix $A \in \mathbb{Z}^{m\times d}$ can be bounded from above by a constant. However, Markov moves can be used twice in a minimal path which can make the diameter of the compressed fiber graph larger than the size of the Markov basis. The next definition puts more constraints on the Markov basis and postulates the existence of a path that uses every move from the Markov basis at most once.

Definition 5.1. Let $\mathcal{F} \subset \mathbb{Z}^d$ be a finite set and $\mathcal{M} = \{m_1, \ldots, m_k\} \subset \mathbb{Z}^d$. An *augmenting path* between distinct $u, v \in \mathcal{F}$ of length $r \in \mathbb{N}$ is a path in $\mathcal{F}^c(\mathcal{M})$ of the form

$$u \to u + \lambda_{i_1} m_{i_1} \to u + \lambda_{i_1} m_{i_1} + \lambda_{i_2} m_{i_2} \to \cdots \to u + \sum_{k=1}^{r} \lambda_{i_k} m_{i_k} = v$$

with distinct indices $i_1, \dots, i_r \in [k]$. An augmenting path is *minimal* for $u, v \in \mathcal{F}$ if there exists no shorter augmenting path between u and v in $\mathcal{F}^c(\mathcal{M})$. A Markov basis $\mathcal{M}$ for $\mathcal{F}$ is *augmenting* if there is an augmenting path between any distinct nodes in $\mathcal{F}$. The *augmentation length* $\mathcal{A}_{\mathcal{M}}(\mathcal{F})$ of an augmenting Markov basis $\mathcal{M}$ is the maximum length of all minimal augmenting paths in $\mathcal{F}^c(\mathcal{M})$.

Not every Markov basis is augmenting (see Example 3.11), but the diameter of compressed fiber graphs that use an augmenting Markov basis is at most the number of the moves. For fiber graphs coming from an integer matrix, an augmenting Markov basis for all of its fibers can be computed (Remark 5.2).

Remark 5.2. Let $A \in \mathbb{Z}^{m\times d}$ with $\ker_{\mathbb{Z}}(A) \cap \mathbb{N}^d = \{0\}$ and let $b \in \mathbb{N}A$. The Graver basis is an augmenting Markov basis for $\mathcal{F}_{A,b}$ for any $b \in \mathbb{N}A$. We claim that when A is totally unimodular, then $\mathcal{A}_{\mathcal{G}_A}(\mathcal{F}_{A,b}) \leq d^2(\operatorname{rank}(A)+1)$. In particular, the augmentation length is independent of the right-hand side b. Let $u, v \in \mathcal{F}_{A,b}$ be arbitrary and for $i \in \mathbb{N}$, let $l_i := \min\{u_i, v_i\}$, $w_i := \max\{u_i, v_i\}$, and $c_i := \operatorname{sign}(u_i - v_i) \in \{-1, 0, 1\}$. Then v is the unique optimal value of the linear integer optimization problem

$$\min\{c^T x : Ax = b, l \leq x \leq w, x \in \mathbb{Z}^d\}.$$

A *discrete steepest decent* as defined in [5, Definition 3] using Graver moves needs at most $\|c\|_1 \cdot d \cdot (\operatorname{rank}(A)+1) \leq d^2 \cdot (\operatorname{rank}(A)+1)$ many augmentations from u to reach the optimal value v. We refer to [5, Corollary 8] which ensures that every Graver move is used at most once. Note that in [5], x is constrained to $x \geq 0$ instead to $x \geq l$, but their argument works for any lower bound.

Example 5.3. Fix $d \in \mathbb{N}$ and consider A and $\mathcal{M}$ from Example 3.5. We show that $\mathcal{M}$ is an augmenting Markov basis for $\mathcal{F}_{A,b}$ for any $b \in \mathbb{N}$. Let $u, v \in \mathcal{F}_{A,b}$ be distinct, then there exists $i \in [d]$ such that $u_i > v_i$ or $u_i < v_i$, thus, we can walk from u to $u' := u + (u_i - v_i)(e_1 - e_i)$ or from v to $v' := v + (v_i - u_i)(e_1 - e_i)$. In any case, after that augmentation, the pairs (u', v) and (v', u) coincide in the ith coordinate and thus we find an augmenting path by induction on the dimension d. We have used at most $d-1$ many edges in these paths and hence $\mathcal{A}_{\mathcal{M}}(\mathcal{F}_{A,b}) \leq d-1$ for all $b \in \mathbb{N}$.

We now show that the augmentation length is essentially bounded from below by the dimension of the node set and hence the bound observed in Example 5.3 cannot be improved. We first need the following lemma.

Lemma 5.4. *Let $v_1, \dots, v_k \in \mathbb{Q}^d$ such that any $v \in \operatorname{span}_{\mathbb{Q}}\{v_1, \dots, v_k\}$ can be represented by a linear combination of r vectors. Then* $\dim(\operatorname{span}_{\mathbb{Q}}\{v_1, \dots, v_k\}) \leq r$.

Proof. Let $\mathfrak{B} \subset \mathfrak{P}(v_1, \dots, v_k)$ the set of all subsets of cardinality r. By our assumption, $\cup_{B\in\mathfrak{B}} \operatorname{span}_{\mathbb{Q}}\{B\} = \operatorname{span}_{\mathbb{Q}}\{v_1, \dots, v_k\}$. Since $\dim(\operatorname{span}_{\mathbb{Q}}\{B\}) \leq r$ for all $B \in \mathfrak{B}$ and since $\mathfrak{B}$ is finite, the claim follows. □

Proposition 5.5. *Let $\mathcal{P} \subset \mathbb{Q}^d$ be polytope and let $\mathcal{M} \subset \mathbb{Z}^d$ be an augmenting Markov basis for $\mathcal{F}_i := (i \cdot \mathcal{P}) \cap \mathbb{Z}^d$ for all $i \in \mathbb{N}$. Then* $\dim(\mathcal{P}) \leq \max_{i\in\mathbb{N}} \mathcal{A}_{\mathcal{M}}(\mathcal{F}_i)$.

Proof. Without restricting generality, we can assume that $0 \in \mathcal{P}$. Let $V := \operatorname{span}_{\mathbb{Q}}\{\mathcal{P}\}$ be the $\mathbb{Q}$-span of $\mathcal{P}$, then $\dim(\mathcal{P}) = \dim(V)$. We must have $\dim(\operatorname{span}_{\mathbb{Q}}\{\mathcal{M}\}) = \dim(V)$ since $\dim(\mathcal{P}) = \dim(\operatorname{conv}_{\mathbb{Q}}(\mathcal{F}_i))$ for i sufficiently large and since $\mathcal{M}$ is a Markov basis for $\mathcal{F}_i$. Define $r := \max_{i\in\mathbb{N}} \mathcal{A}_{\mathcal{M}}(\mathcal{F}_i)$ and choose any non-zero $v \in V$ and $u \in \operatorname{relint}(\mathcal{P}) \subset \mathbb{Q}^d$. Then there exists $\delta \in \mathbb{Q}_{>0}$ such that $u + \delta v \in \mathcal{P}$. Thus, $\frac{1}{\delta}u + v \in \frac{1}{\delta}\mathcal{P}$. Let $c \in \mathbb{N}_{\geq 1}$ such that $i := \frac{c}{\delta} \in \mathbb{N}$ and $w := \frac{c}{\delta}u \in \mathbb{Z}^d$. Then $w + cv = c(\frac{1}{\delta}u + v) \in (i \cdot \mathcal{P}) \cap \mathbb{Z}^d = \mathcal{F}_i$. By assumption, there exists an augmenting path from w to $w + cv$ using only r elements from $\mathcal{M}$. Put differently, the element cv from V can be represented by a linear combination of r vectors from $\mathcal{M}$. Since v was chosen arbitrarily, Lemma 5.4 implies $\dim(\mathcal{P}) = \dim(V) \leq r$. □

Remark 5.6. It is a consequence from Proposition 5.5 that for any matrix $A \in \mathbb{Z}^{m\times d}$ with $\ker_{\mathbb{Z}}(A) \cap \mathbb{N}^d = \{0\}$ and an augmenting Markov basis $\mathcal{M}$, there exists $\mathcal{F} \in \mathcal{P}_A$ such that $\mathcal{A}_{\mathcal{M}}(\mathcal{F}) \geq \dim(\ker_{\mathbb{Z}}(A))$.

Let us now shortly recall the framework from [23] which is necessary to prove our main theorem. Let $G = (V, E)$ be a graph. For any ordered pair of distinct nodes $(x, y) \in V \times V$, let $p_{x,y} \subseteq E$ be a path from x to y in G and let $\Gamma := \{p_{x,y} : (x, y) \in V \times V, x \neq y\}$ be the collection of these paths, then Γ is *a set of canonical paths*. Let for any edge $e \in E$, $\Gamma_e := \{p \in \Gamma : e \in p\}$ be the set of paths from Γ that use e. Now, let $\mathcal{H} : V \times V \to [0, 1]$ be a symmetric random walk on G and define

$$\rho(\Gamma, \mathcal{H}) := \frac{\max\{|p| : p \in \Gamma\}}{|V|} \cdot \max_{\{u,v\}\in E} \frac{|\Gamma_{\{u,v\}}|}{\mathcal{H}(u, v)}.$$

Observe that symmetry of $\mathcal{H}$ is needed to make $\rho(\Gamma, \mathcal{H})$ well-defined. This can be used to prove the following upper bound on the second largest eigenvalue.

Lemma 5.7. *Let G be a graph, $\mathcal{H}$ be a symmetric random walk on G, and Γ be a set of canonical paths in G. Then $\lambda_2(\mathcal{H}) \leq 1 - \frac{1}{\rho(\Gamma,\mathcal{H})}$.*

Proof. The stationary distribution of $\mathcal{H}$ is the uniform distribution and thus the statement is a direct consequence of [23, Theorem 5], since $\rho(\Gamma, \mathcal{H})$ is an upper bound on the constant defined in [23, equation 4]. □

Theorem 5.8. *Let $\mathcal{F} \subset \mathbb{Z}^d$ be finite and let $\mathcal{M} := \{m_1, \ldots, m_k\} \subset \mathbb{Z}^d$ be an augmenting Markov basis. Let π be the uniform and f be a positive distribution on $\mathcal{F}$ and $\mathcal{M}$ respectively. For $i \in [k]$, let $r_i := \max\{|\mathcal{R}_{\mathcal{F},m_i}(u)| : u \in \mathcal{F}\}$ and suppose that $r_1 \geq r_2 \geq \cdots \geq r_k$. Then*

$$\lambda(\mathcal{H}^{\pi,f}_{\mathcal{M},\mathcal{F}}) \leq 1 - \frac{|\mathcal{F}| \cdot \min(f)}{\mathcal{A}_{\mathcal{M}}(\mathcal{F}) \cdot \mathcal{A}_{\mathcal{M}}(\mathcal{F})! \cdot 3^{\mathcal{A}_{\mathcal{M}}(\mathcal{F})-1} \cdot 2^{|\mathcal{M}|} \cdot r_1 r_2 \cdots r_{\mathcal{A}_{\mathcal{M}}(\mathcal{F})}}.$$

Proof. Choose for any distinct $u, v \in \mathcal{F}$ an augmenting path $p_{u,v}$ of minimal length in $\mathcal{F}^c(\mathcal{M})$ and let Γ be the collection of all these paths. Let $u + \mu m_k = v$ be an edge in $\mathcal{F}^c(\mathcal{M})$, then our goal is to bound $|\Gamma_{\{u,v\}}|$ from above. Let $\mathcal{S} := \{S \subseteq [r] : |S| \leq \mathcal{A}_{\mathcal{M}}(\mathcal{F}), k \in S\}$ and take any path $p_{x,y} \in \Gamma_{\{u,v\}}$. Then there exists $S := \{i_1, \ldots, i_s\}$ with $s := |S| \leq \mathcal{A}_{\mathcal{M}}(\mathcal{F})$ such that $x + \sum_{k=1}^{s} \lambda_{i_k} m_{i_k} = y$. Since $p_{x,y}$ uses the edge $\{u, v\}$, there is $j \in [s]$ such that $i_j = k$ and $\lambda_{i_j} = \mu$. Since $|\lambda_{i_k}| \leq r_{i_k}$, there are at most

$$s! \cdot (2r_{i_1} + 1) \cdots (2r_{i_{j-1}} + 1) \cdot (2r_{i_{j+1}} + 1) \cdots (2r_{i_s} + 1) \leq s! \cdot 3^{s-1} \prod_{t\in S\setminus\{k\}} r_t$$

paths in $\Gamma_{\{u,v\}}$ that uses the edge $\{u,v\}$ and the moves $m_{i_1},\dots,m_{i_{j-1}},m_{i_{j+1}}\dots,m_{i_s}$. Since all the paths are minimal, they have length at most $\mathcal{A}_{\mathcal{M}}(\mathcal{F})$ so indeed every path in Γ has that form.

$$\frac{|\Gamma_{u,v}|}{\mathcal{H}^{\pi,f}_{\mathcal{F},\mathcal{M}}(u,v)} \leq 3^{\mathcal{A}_{\mathcal{M}}(\mathcal{F})-1}\frac{\sum_{S\in\mathcal{S}}\left(|S|!\prod_{t\in S\setminus\{k\}} r_t\right)}{f(m_{i_j})\cdot\frac{1}{|\mathcal{R}_{m_{i_j}}(u)|}} \leq \frac{3^{\mathcal{A}_{\mathcal{M}}(\mathcal{F})-1}\cdot\mathcal{A}_{\mathcal{M}}(\mathcal{F})!\cdot|\mathcal{S}|\cdot r_1r_2\dots r_{\mathcal{A}_{\mathcal{M}}(\mathcal{F})}}{f(m_{i_j})},$$

where we have used the assumption $r_1\geq r_2\geq\cdots\geq r_k$. Bounding $|\mathcal{S}|$ rigorously from above by $2^{|\mathcal{M}|}$, the claim follows from Lemma 5.7. □

Definition 5.9. Let $\mathcal{F}\subset\mathbb{Z}^d$ and $\mathcal{M}\subset\mathbb{Z}^d$ be finite sets. The longest ray through $\mathcal{F}$ along vectors of $\mathcal{M}$ is $\mathcal{R}_{\mathcal{F},\mathcal{M}} := \arg\max\{|\mathcal{R}_{\mathcal{F},m}(u)| : m\in\mathcal{M}, u\in\mathcal{F}\}$.

Corollary 5.10. *Let $(\mathcal{F}_i)_{i\in\mathbb{N}}$ be a sequence of finite sets in $\mathbb{Z}^d$ and let π_i be the uniform distribution on $\mathcal{F}_i$. Let $\mathcal{M}\subset\mathbb{Z}^d$ be an augmenting Markov basis for $\mathcal{F}_i$ with $\mathcal{A}_{\mathcal{M}}(\mathcal{F}_i)\leq\dim(\mathcal{F}_i)$ and suppose that $(|\mathcal{R}_{\mathcal{F}_i,\mathcal{M}}|)^{\dim(\mathcal{F}_i)})_{i\in\mathbb{N}}\in\mathcal{O}(|\mathcal{F}_i|)_{i\in\mathbb{N}}$. Then for any positive mass function $f:\mathcal{M}\to[0,1]$, there exists $\epsilon>0$ such that $\lambda(\mathcal{H}^{\pi_i,f}_{\mathcal{F}_i,\mathcal{M}})\leq 1-\epsilon$ for all $i\in\mathbb{N}$.*

Proof. This is a straightforward application of Theorem 5.8. □

Corollary 5.11. *Let $\mathcal{P}\subset\mathbb{Z}^d$ be a polytope, $\mathcal{F}_i := (i\cdot\mathcal{P})\cap\mathbb{Z}^d$ for $i\in\mathbb{N}$, and let π_i be the uniform distribution on $\mathcal{F}_i$. Suppose that $\mathcal{M}\subset\mathbb{Z}^d$ is an augmenting Markov basis $\{\mathcal{F}_i : i\in\mathbb{N}\}$ such that $\mathcal{A}_{\mathcal{M}}(\mathcal{F}_i)\leq\dim(\mathcal{P})$ for all $i\in\mathbb{N}$. Then for any positive mass function $f:\mathcal{M}\to[0,1]$, there exists $\epsilon>0$ such that $\lambda(\mathcal{H}^{\pi_i,f}_{\mathcal{F}_i,\mathcal{M}})\leq 1-\epsilon$ for all $i\in\mathbb{N}$.*

Proof. Let $r:=\dim(\mathcal{P})$. We first show that $(|\mathcal{R}_{\mathcal{F}_i,\mathcal{M}}|)_{i\in\mathbb{N}}\in\mathcal{O}(i)_{i\in\mathbb{N}}$. Write $\mathcal{M}=\{m_1,\dots,m_k\}$ and denote by $l_i := \max\{|(u+m_i\cdot\mathbb{Z})\cap\mathcal{P}| : u\in\mathcal{P}\}$ be the length of the longest ray through the polytope $\mathcal{P}$ along m_i. It suffices to prove that $i\cdot(l_k+1)$ is an upper bound on the length of any ray along m_k through $\mathcal{F}_i$. For that, let $u\in\mathcal{F}_i$ such that $u+\lambda m_k\in\mathcal{F}_i$ for some $\lambda\in\mathbb{N}$, then $\frac{1}{i}u+\frac{\lambda}{i}m_k\in\mathcal{P}$ and thus $\lfloor\frac{\lambda}{i}\rfloor\leq l_k$, which gives $\lambda\leq i\cdot(l_k+1)$. With $C := \max\{l_1,\dots,l_k\}+1$ we have $|\mathcal{R}_{\mathcal{F}_i,\mathcal{M}}|\leq C\cdot i$. Ehrhart's theorem [2, Theorem 3.23] gives $(|\mathcal{F}_i|)_{i\in\mathbb{N}}\in\Omega(i^r)_{i\in\mathbb{N}}$ and since $|\mathcal{R}_{\mathcal{F}_i,\mathcal{M}}|\leq C\cdot i$, we have $(|\mathcal{R}_{\mathcal{F}_i,\mathcal{M}}|^r)_{i\in\mathbb{N}}\in\mathcal{O}(|\mathcal{F}_i|)_{i\in\mathbb{N}}$. An application of Corollary 5.10 proves the claim. □

Example 5.12. Fix $d,r\in\mathbb{N}$ and let $\mathcal{C}_{d,r} := \{u\in\mathbb{Z}^d : \|u\|_1\leq r\}$ be the set of integers of the d-dimensional cross-polytope with radius r. The set $\mathcal{M}_d=\{e_1,\dots,e_d\}$ is a Markov basis for $\mathcal{C}_{d,r}$ for any $r\in\mathbb{N}$. We show that $\mathcal{M}_d$ is an augmenting Markov basis whose augmentation length is at most d. For that, let $u,v\in\mathcal{C}_{d,r}$ distinct elements. We claim that there exists $i\in[d]$ such that $x_i\neq v_i$ and $u_i+(v_i-u_i)\in\mathcal{C}_{d,r}$. Let $S\subseteq[d]$ be the set of indices where u and v differ and let $s=r-||u||_1$. If $|S|=1$, then the result is clear so suppose $|S|\geq 2$. If the result doesn't hold then for all $i\in S$, $|v_i|-|u_i|>s$. It follows that

$$\|v\|_1=\sum_{i\notin S}|u_i|+\sum_{i\in S}|v_i|>\sum_{i\notin S}|u_i|+\sum_{i\in S}s+|u_i|=|S_{uv}|\cdot s+\|u\|_1=(|S|-1)\cdot s+r.$$

But we assumed that $v\in\mathcal{C}_{d,r}$. It follows that for any pair of points u,v in $\mathcal{C}_{d,r}$, there is a walk, using the unit vectors as moves, that uses each move at most once. Corollary 5.10 yield that for any $d\in\mathbb{N}$, the second largest eigenvalue modulus of the heat-bath random walk on $\mathcal{C}_{d,r}$ with uniform as stationary distribution can be strictly bounded away from 1 for $r\to\infty$.

The bound on the second largest eigenvalue in Theorem 5.8 is quite general and can be improved vastly, provided one has better control over the paths. For example, this can be achieved for hyperrectangles intersected with a halfspace.

Proposition 5.13. *Let $a \in \mathbb{N}_{>0}^d$, $b \in \mathbb{N}$, $\mathcal{F} = \{u \in \mathbb{N}^d : a^T \cdot u \leq b\}$, and $\mathcal{M} := \{e_1, \dots, e_d\}$. If π and f are the uniform distributions on $\mathcal{F}$ and $\mathcal{M}$ respectively, then*

$$\lambda(\mathcal{H}^{\pi,f}_{\mathcal{F},\mathcal{M}}) \leq 1 - \frac{|\mathcal{F}|}{d^2} \prod_{i=1}^{d} \frac{a_i}{b}.$$

Proof. Observe that $\mathcal{M}$ is a Markov basis for $\mathcal{F}$ since all nodes are connected with $0 \in \mathcal{F}$. Let $u, v \in \mathcal{F}$ be distinct. We first show that there exists $k \in [d]$ such that $u_k \neq v_k$ and $u + (v_k - u_k)e_k \in \mathcal{F}$. If $u \leq v$, the statement trivially holds. Otherwise, there exists $k \in [d]$ such that $u_k > v_k$ and the vector obtained by replacing the kth coordinate of u by v_k remains in $\mathcal{F}$. Now, consider for the following path between u and v: Choose the smallest index $k \in [d]$ such that $u_k \neq v_k$ and such that $u + (v_k - u_k) \cdot e_k \in \mathcal{F}$ and proceed recursively with $u + (v_k - u_k)$ and v. This gives a path $p_{u,v}$ between u and v of length at most d. Let Γ be the collection of all these paths. We want to apply Lemma 5.7. Thus, let $x \in \mathcal{F}$ and consider the edge $x \to x + c \cdot e_s$. Let us count the paths $p_{u,v}$ that use that edge. Let $u, v \in \mathcal{F}$ and let $k_1, \dots, k_r \in [d]$ be distinct indices such that

$$u \to u + (v_{k_1} - u_{k_1})e_{k_1} \to u + (v_{k_1} - u_{k_1})e_{k_1} + (v_{k_2} - u_{k_2})e_{k_2} \to \cdots \to v$$

represents the path $p_{u,v}$ constructed by the upper rule. Assume that $p_{u,v}$ uses the edge $\{x, x + ce_s\}$ and let $k_l = s$ and $(v_{k_l} - u_{k_l}) = c$. In particular,

$$u + (v_{k_1} - u_{k_1})e_{k_1} + \cdots + (v_{k_{l-1}} - u_{k_{l-1}})e_{k_{l-1}} = x$$
$$x + (v_{k_l} - u_{k_l})e_{k_l} + \cdots + (v_{k_r} - u_{k_r})e_{k_r} = v.$$

We see that $v_{k_t} = x_{k_t}$ for all $t < l$ and that $u_{k_t} = x_{k_t}$ for all $t \geq l$. In particular, $v_{k_l} = u_{k_l} + c = x_{k_l} + c$ is also fixed. The coordinates u_{k_t} and v_{k_t} are bounded from above by $\frac{b}{a_{k_t}}$ for all $t \in [r]$, and hence there can be at most

$$\left(\prod_{t=1}^{l-1} \frac{b}{a_{k_t}}\right) \cdot \left(\prod_{t=l+1}^{r} \frac{b}{a_{k_t}}\right).$$

Since $k_1, \dots, k_t$ are distinct coordinate indices, we have

$$\frac{|\Gamma_{x,x+c\cdot e_s}|}{\mathcal{H}^{\pi,f}_{\mathcal{F},\mathcal{M}}(x, x + c \cdot e_s)} \leq d \cdot \prod_{i=1}^{d} \frac{b}{a_i}.$$

Lemma 5.7 finishes the proof. □

In fixed dimension, Proposition 5.13 leads to rapid mixing, but for $d \to \infty$, no statement can be made. In [19], it was shown that the simple walk with an additional halting probability on $\{u \in \mathbb{N}^d : a^t u \leq b\} \cap \{0, 1\}^d$ has mixing time in $\mathcal{O}(d^{4.5+\epsilon})$. For zero-one polytopes, simple and heat-bath walk coincide and we are confident that a similar statement holds without the restriction on zero-one polytopes.

The heat-bath random walk mixes rapidly when an augmenting Markov basis with a small augmentation length is used. We think that it is interesting to question how might an augmenting Markov bases be obtained and how their augmentation length can be improved.

Question 5.14. *Let $\mathcal{M}$ be an augmenting Markov basis of A. Can we find finitely many moves $m_1, \ldots, m_k$ such that the augmentation length of $\mathcal{M} \cup \{m_1, \ldots, m_k\}$ on $\mathcal{F}_{A,b}$ is at most* $\dim(\ker_{\mathbb{Z}}(A))$ *for all $b \in \mathbb{N}A$?*

References

1. Stephen Baumert, Archis Ghate, Seksan Kiatsupaibul, Yanfang Shen, Robert L. Smith, and Zelda B. Zabinsky, *Discrete Hit-and-Run for Sampling Points from Arbitrary Distributions Over Subsets of Integer Hyperrectangles*, Operations Research **57** (2009), no. 3, 727–739.
2. Matthias Beck and Sinai Robins, *Computing the Continuous Discretely*, Springer, New York, 2007.
3. Mary Cryan, Martin Dyer, Leslie Ann Goldberg, Mark Jerrum, and Russell Martin, *Rapidly Mixing Markov Chains for Sampling Contingency Tables with a Constant Number of Rows*, SIAM Journal on Computing **36** (2006), no. 1, 247–278.
4. Jesús A. De Loera, Raymond Hemmecke, and Matthias Köppe, *Algebraic and Geometric Ideas in the Theory of Discrete Optimization*, MPS-SIAM Series on Optimization, SIAM, Cambridge, 2013.
5. Jesús A. De Loera, Raymond Hemmecke, and Jon Lee, *On Augmentation Algorithms for Linear and Integer-Linear Programming: From Edmonds–Karp to Bland and Beyond*, SIAM Journal on Optimization **25** (2015), no. 4, 2494–2511.
6. Persi Diaconis and Bernd Sturmfels, *Algebraic algorithms for sampling from conditional distributions*, The Annals of statistics **26** (1998), no. 1, 363–397.
7. Mathias Drton, Bernd Sturmfels, and Seth Sullivant, *Lectures on algebraic statistics*, Oberwolfach Seminars, vol. 39, Springer, Berlin, 2009, A Birkhäuser book.
8. Martin Dyer, Catherine Greenhill, and Mario Ullrich, *Structure and eigenvalues of heat-bath Markov chains*, Linear Algebra and its Applications **454** (2014), 57–71.
9. Giles Gardam, *Expander Graphs and Kazhdan' s Property (T)*, Bachelor's thesis, University of Sidney, 2012.
10. Hisayuki Hara, Akimichi Takemura, and Ruriko Yoshida, *On connectivity of fibers with positive marginals in multiple logistic regression*, Journal of Multivariate Analysis (2010), 1–26.
11. Raymond Hemmecke and Peter N. Malkin, *Computing generating sets of lattice ideals and Markov bases of lattices*, Journal of Symbolic Computation **44** (2009), no. 10, 1463–1476.
12. Roger A. Horn, *Matrix Analysis*, 2nd ed., Cambridge University Press, New York, 2013.
13. Steven S. Kim, *Mixing Time of a Rook's Walk*, Undergraduate certificate paper (2012).
14. David A. Levin, Yuval Peres, and Elisabeth L. Wilmer, *Markov chains and mixing times*, American Mathematical Society, Providence, RI, 2009.
15. László Lovász, *Hit-and-run mixes fast*, Mathematical Programming **86** (1999), no. 3, 443–461.
16. László Lovász and Santosh Vempala, *Hit-and-Run from a Corner*, SIAM Journal on Computing **35** (2006), no. 4, 985–1005.
17. Peter N. Malkin, *Computing Markov bases, Gröbner bases, and extreme rays*, Phd thesis, 2007, p. 223.
18. Cam McLeman, Peter T. Otto, John Rahmani, and Matthew Sutter, *Mixing times for the Rook's walk via path coupling*, to appear in Involve (2016), 1–12.
19. Ben Morris and Alistair Sinclair, *Random Walks on Truncated Cubes and Sampling 0-1 Knapsack Solutions*, SIAM Journal on Computing **34** (2004), no. 1, 195–226.
20. Samu Potka, *Higher connectivity of fiber graphs of Gröbner bases*, Journal of Algebraic Statistics **4** (2013), no. 1, 93–107.
21. Johannes Rauh and Seth Sullivant, *Lifting Markov bases and higher codimension toric fiber products*, Journal of Symbolic Computation **74** (2016), 276–307.
22. András Sebö, *Hilbert Bases, Caratheodory's Theorem and Combinatorial Optimization*, (1990), 431–455.

23. Alistair Sinclair, *Improved Bounds for Mixing Rates of Markov Chains and Multicommodity Flow*, Combinatorics, Probability and Computing **1** (1992), no. 4, 351–370.
24. Bernd Sturmfels, *Gröbner bases and convex polytopes*, American Mathematical Society, Providence, R.I., 1996.
25. Seth Sullivant, *Markov bases of binary graph models*, Annals of Combinatorics **7** (2003), 441–466.
26. Santosh S. Vempala, *Geometric Random Walks: A Survey*, MSRI Combinatorial and Computational Geometry **52** (2005), 573–612.
27. Tobias Windisch, *Rapid mixing and Markov bases*, preprint, arXiv:1505.03018 (2015), 1–18.

NC State University, Raleigh, NC 27695, USA
E-mail address: crstanl2@ncsu.edu

Otto-von-Guericke Universität, Magdeburg, Germany
E-mail address: windisch@ovgu.de